



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/674,942	09/29/2003	Paul A. Martin	SUN-P9580	5182

57960 7590 02/05/2007
SUN MICROSYSTEMS INC.
C/O PARK, VAUGHAN & FLEMING LLP
2820 FIFTH STREET
DAVIS, CA 95618-7759

EXAMINER

RUTZ, JARED IAN

ART UNIT	PAPER NUMBER
----------	--------------

2187

SHORTENED STATUTORY PERIOD OF RESPONSE	MAIL DATE	DELIVERY MODE
3 MONTHS	02/05/2007	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

If NO period for reply is specified above, the maximum statutory period will apply and will expire 6 MONTHS from the mailing date of this communication.

Office Action Summary

Application No.

10/674,942

Applicant(s)

MARTIN ET AL.

Examiner

Jared I. Rutz

Art Unit

2187

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 18 September 2006.
- 2a) ☐ This action is **FINAL**. 2b) ☒ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1-7, 9-11, 13-21, 23-25, 27-35, 37-39, 41 and 42 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 1-7, 9-11, 13-21, 23-25, 27-35, 37-39, and 41-42 is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☐ The drawing(s) filed on _____ is/are: a) ☐ accepted or b) ☐ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some * c) ☐ None of:
- ☐ Certified copies of the priority documents have been received.
 - ☐ Certified copies of the priority documents have been received in Application No. _____.
 - ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).
- * See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- ☐ Notice of References Cited (PTO-892)
- ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
- ☐ Information Disclosure Statement(s) (PTO/SB/08)
Paper No(s)/Mail Date _____.
- ☐ Interview Summary (PTO-413)
Paper No(s)/Mail Date. _____.
- ☐ Notice of Informal Patent Application
- ☐ Other: _____.

DETAILED ACTION

1. Claims 1-7, 9-11, 13-21, 23-25, 27-35, 37-39, and 41-42 are pending in the instant application. The amendments filed 9/18/2006 were entered with the filing of a Request for Continued Examination on 11/20/2006. Applicant's arguments filed 9/18/2006 have been fully and carefully considered, but are not found persuasive.

Claim Rejections - 35 USC § 101

2. 35 U.S.C. 101 reads as follows:

Whoever invents or discovers any new and useful process, machine, manufacture, or composition of matter, or any new and useful improvement thereof, may obtain a patent therefor, subject to the conditions and requirements of this title.

3. **Claims 15-21, 23-25, and 27-28** are rejected under 35 U.S.C. 101 because the claimed invention is directed to non-statutory subject matter. Claim 15 is directed to a computer-readable storage medium. Paragraph 0041 states "*...a computer-readable storage medium, which may be any device or medium that can store code and/or data for use by a computer system. This includes, but is not limited to, magnetic and optical storage devices such as disk drives, magnetic tape, CDs (compact discs) and DVDs (digital versatile discs or digital video disks), and computer instruction signals embodied in a transmission medium (with or without a carrier wave upon which the signals are modulated). For example, the transmission medium may include a communications network, such as the Internet.*" The claimed computer-readable storage medium is not limited to a tangible product, but may merely comprise an energy field, which is not a process, machine, manufacture, or a new or useful improvement thereof. Additionally, a transmission medium is not capable of storing computer instruction signals, but provides

Art Unit: 2187

a path by which an signal (electromagnetic, optical, or the like) travels with the encoded computer instructions. Therefore, a transmission medium, which may not be limited to a tangible embodiment, is not capable of performing the claimed limitations. Accordingly, claim 15 is not limited to statutory subject matter.

4. **Claims 16-21, 23-25, and 27-28** depend from claim 15, and are also directed to a computer-readable storage medium. Said claims are rejected for the same reasons as claim 15.

5. To overcome this rejection, applicant is advised to either cancel the rejected claims, or disavow the inclusion of computer instruction signals embodied in a transmission medium from the scope of a computer-readable storage medium.

Claim Rejections - 35 USC § 103

6. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

7. **Claims 1-3, 9-11, 13, 15-17, 23-25, 27, 29-31 and 37-39 and 41** are rejected under 35 U.S.C. 103(a) as being unpatentable over Michael ("High Performance Dynamic Lock-Free Hash Tables and List-Based Sets") in view of applicants' admitted prior art (AAPA), and further in view of Schimmel (US 5,960,434).

Art Unit: 2187

8. With respect to **claims 1, 15 and 29**, Michael discloses a method (which can also be implemented in a computer-readable storage medium, or by an apparatus) for using a hash table that is fully dynamic and lock-free (page 73, column 2, paragraph 7, "This paper presents..."), comprising:

- a. performing a lookup into the hash table, wherein the lookup involves using a hash key to lookup a bucket pointer in a bucket array (page 74, column 1, paragraph 5, lines 3-8, "The most common method..."),
- b. following the bucket pointer to a data node within a linked list (page 74, column 1, paragraph 5, lines 3-8, "The most common method..."; page 77, column 2, paragraph 5, "A thread keeps..."), and
- c. searching from the data node through the linked list to locate a node that matches the hash key if one exists (page 77, column 2, paragraph 5, "A thread keeps...")
- d. wherein the linked list contains only data nodes and at most a constant number of dummy nodes (page 74, Fig. 1 shows only data nodes; page 74, column 1, paragraph 10, "Figure 1 shows...").

9. Michael does not disclose the limitations wherein

- e. the linked list contains all of the data nodes in the hash table;
- f. if the average number of data nodes in each bucket exceeds a maximum value, the method further comprises:
- g. increasing the number of buckets in the bucket array to form a larger bucket array, and

h. using more bits from the hash key to perform lookups in the larger bucket array,

i. wherein the data nodes are stored in the linked list in bit-inverted hash key order, and

j. wherein increasing the number of buckets in the bucket array involves copying the existing bucket array into the top half of the larger bucket array.

10. However, AAPA disclose the limitation wherein the linked list contains all of the data nodes in the hash table (Fig. 2A; specification, page 4, lines 3-6).

11. Michael and AAPA are analogous art because they are from the same field of endeavor, namely hash table design.

12. At the time of invention, it would have been obvious to a person of ordinary skill in the art to combine the single linked list that contains all the data nodes of the hash table of the AAPA with the hash table of Michael.

13. As noted by applicant's "Related Art" in the specification on page 4, lines 5-6, the motivation for doing so would have been because having a single linked list that contains all of the data nodes prevents the data nodes from "needing to be moved when the number of hash buckets changes".

14. Additionally, Schimmel discloses the limitations wherein if the average number of data nodes in each bucket exceeds a maximum value, the method further comprises:

k. increasing the number of buckets in the bucket array to form a larger bucket array (column 2, lines 29-41); and

l. using more bits from the hash key to perform lookups in the larger bucket array (column 5, lines 2-8 say that the hash keys are represented by the numbers in the individual data records in Figs. 1A and 1B, and the keys range from 0-15, which may be represented by four bits; Fig. 11 shows a table with more buckets including a hash key with a value of 28, which requires a fifth bit to be represented in binary).

15. Michael, AAPA and Schimmel are analogous art because they are from the same field of endeavor, namely hash table design.

16. At the time of invention, it would have been obvious to a person of ordinary skill in the art to combine the dynamic addition of hash table buckets of Schimmel with the hash table of Michael and the AAPA.

17. The motivation for doing so would have been because this system is able to "keep [the hash table structure] as compact as possible and to avoid long bucket links" (Schimmel, column 2, line 25-26).

18. Schimmel also discloses the limitations

m. wherein the data nodes are stored in the linked list in bit-inverted hash key order (in Figs. 1A, 3A, 4A and 4B, the nodes are stored in decreasing bit-inverted hash key order); and

n. wherein increasing the number of buckets in the bucket array involves copying the existing bucket array into the top half of the larger bucket array (Figures 3A and 4A, column 6 lines 64-67).

Art Unit: 2187

19. At the time of invention, it would have been obvious to a person of ordinary skill in the art to combine the storing of data nodes in bit-inverted hash key order and the mapping of the existing bucket array into the top half of the larger bucket array of Schimmel with the hash table of Michael and the AAPA.

20. The motivation for doing so would have been because Schimmel teaches that "re-hashing of data records 156 in step 818 can take too long to be practical" (column 7, lines 39-41).

21. Therefore, it would have been obvious to a person of ordinary skill in the art to combine the AAPA and Schimmel with Michael for the benefit of a hash table with a single linked list that contains all the data nodes table, and that dynamically adds hash buckets onto itself, stores data nodes in bit-inverted hash key order and maps the existing bucket array into the top half of the larger bucket array to obtain the invention as specified in claims 1, 15 and 29.

22. With respect to **claims 2, 16 and 30**, Michael in view of the AAPA and Schimmel disclose the method (which can also be implemented in a computer-readable storage medium, or by an apparatus) of claims 1, 15 and 29 as shown supra. Michael further discloses the limitation wherein the data node pointed to by the bucket pointer precedes the nodes in the bucket (Fig. 1; page 74, column 1, paragraph 10).

23. With respect to **claims 3, 17 and 31**, Michael in view of the AAPA and Schimmel disclose the method (which can also be implemented in a computer-readable storage

medium, or by an apparatus) of claims 1, 15 and 29 as shown supra. Michael further discloses the limitation wherein deleting the data node from the linked list involves:

- o. using an atomic operation to mark the data node as dead (page 76, column 2, paragraph 6, lines 3-5, "If the key is found..."); and
- p. atomically updating the next pointer of the predecessor of the data node to point around the data node to the successor of the data node in the linked list (Fig. 6; page 76, column 2, paragraph 6, "A Delete operation...").

24. With respect to **claims 9, 23 and 37**, Michael in view of the AAPA and Schimmel disclose the method (which can also be implemented in a computer-readable storage medium, or by an apparatus) of claims 1, 15 and 29 as shown supra. Michael in view of the AAPA do not disclose the limitation wherein buckets in the larger bucket array are initialized on-the-fly as they are referenced.

25. However, Schimmel discloses the limitation wherein buckets in the larger bucket array are initialized on-the-fly as they are referenced (column 7, lines 47-50; column 8, lines 9-22).

26. Michael, AAPA and Schimmel are analogous art because they are from the same field of endeavor, namely hash table design.

27. At the time of invention, it would have been obvious to a person of ordinary skill in the art to combine the on-the-fly initialization of hash table buckets of Schimmel with the hash table of Michael and the AAPA.

Art Unit: 2187

28. The motivation for doing so would have been because "re-hashing of data records 156 in step 818 can take too long to be practical" (column 7, lines 39-41).

29. Therefore, it would have been obvious to a person of ordinary skill in the art to combine Schimmel with Michael and the AAPA for the benefit of a hash table that initializes new bucket arrays on-the-fly to obtain the invention as specified in claims 9, 23 and 37.

30. With respect to **claims 10, 24 and 38**, Michael in view of the AAPA and Schimmel disclose the method (which can also be implemented in a computer-readable storage medium, or by an apparatus) of claims 1, 15 and 29 as shown supra. Michael in view of the AAPA do not disclose the limitations wherein initializing a bucket pointer involves:

- q. obtaining a parent bucket pointer for the bucket pointer;
- r. searching through the linked list from a node pointed to by the parent bucket pointer to locate a starting node for the bucket pointer; and
- s. updating the bucket pointer to point to the starting node.

31. However, Schimmel discloses the limitations wherein initializing a bucket pointer involves:

- t. obtaining a parent bucket pointer for the bucket pointer (column 8, lines 11-16; "if a bucket... includes a logical back pointer," lines 13-14, indicate that the bucket pointer is being initialized);

Art Unit: 2187

- u. searching through the linked list from a node pointed to by the parent bucket pointer to locate a starting node for the bucket pointer (column 8, lines 16-20); and
- v. updating the bucket pointer to point to the starting node (column 8, lines 20-22).

32. Michael, AAPA and Schimmel are analogous art because they are from the same field of endeavor, namely hash table design.

33. At the time of invention, it would have been obvious to a person of ordinary skill in the art to combine the initialization of bucket pointers on-the-fly of Schimmel with the hash table of Michael and the AAPA.

34. The motivation for doing so would have been because "re-hashing of data records 156 in step 818 can take too long to be practical" (column 7, lines 39-41).

35. Therefore, it would have been obvious to a person of ordinary skill in the art to combine Schimmel with Michael and the AAPA for the benefit of a hash table that initializes new bucket pointers on-the-fly to obtain the invention as specified in claims 10, 24 and 38.

36. With respect to **claims 11, 25 and 39**, Michael in view of the AAPA discloses the method (which can also be implemented in a computer-readable storage medium, or by an apparatus) of claims 1, 15 and 29 as shown supra. Michael in view of the AAPA do not disclose the limitations wherein, if there exists an old hash table, initializing a bucket

Art Unit: 2187

pointer involves looking for a corresponding entry in the old hash table first, and if this fails:

- w. obtaining a parent bucket pointer for the bucket pointer;
- x. searching through the linked list from a node pointed to by the parent bucket pointer to locate a starting node for the bucket pointer; and
- y. updating the bucket pointer to point to the starting node.

37. However, Schimmel discloses the limitations wherein if there exists an old hash table, initializing a bucket-pointer involves looking for a corresponding entry in the old hash table first (column 8, lines 13-14, if the bucket contains a regular pointer), and if this fails:

- z. obtaining a parent bucket pointer for the bucket pointer (column 8, lines 11-16; "if a bucket... includes a logical back pointer," lines 13-14, indicate that the bucket pointer is being initialized);
- aa. searching through the linked list from a node pointed to by the parent bucket pointer to locate a starting node for the bucket pointer (column 8, lines 16-20); and
- bb. updating the bucket pointer to point to the starting node (column 8, lines 20-22).

38. Michael, AAPA and Schimmel are analogous art because they are from the same field of endeavor, namely hash table design.

Art Unit: 2187

39. At the time of invention, it would have been obvious to a person of ordinary skill in the art to combine the initialization of bucket pointers on-the-fly of Schimmel with the hash table of Michael and the AAPA.

40. The motivation for doing so would have been because "re-hashing of data records 156 in step 818 can take too long to be practical" (column 7, lines 39-41).

41. Therefore, it would have been obvious to a person of ordinary skill in the art to combine Schimmel with Michael and the AAPA for the benefit of a hash table that initializes new bucket pointers on-the-fly to obtain the invention as specified in claims 11, 25 and 39.

42. With respect to **claims 13, 27 and 41**, Michael in view of the AAPA discloses the method (which can also be implemented in a computer-readable storage medium, or by an apparatus) of claims 1, 15 and 29 as shown supra. Michael in view of the AAPA do not disclose the limitations

cc. wherein the data nodes are stored in the linked list in hash key order; and

dd. wherein increasing the number of buckets in the bucket array involves interleaving the bucket array into the larger bucket array.

43. However, Schimmel discloses the limitations

ee. wherein the data nodes are stored in the linked list in hash key order (Figs. 1A, 3A, 4A and 4B, where the hash keys are the numbers in each node);
and

ff. wherein increasing the number of buckets in the bucket array involves interleaving the bucket array into the larger bucket array (as the data nodes are rehashed from Fig. 4A to Fig. 4B, they are interleaved among the old and new buckets).

44. Michael, AAPA and Schimmel are analogous art because they are from the same field of endeavor, namely hash table design.

45. At the time of invention, it would have been obvious to a person of ordinary skill in the art to combine the storing of nodes in hash key order and the interleaving of data nodes in the expanded buckets of Schimmel with the hash table of Michael and the AAPA.

46. The motivation for doing so would have been because "only half of the data records need to be moved" after extending the number of buckets (column 7, lines 22-23).

47. Therefore, it would have been obvious to a person of ordinary skill in the art to combine Schimmel with Michael and the AAPA for the benefit of a hash table that stores data nodes in hash key order and interleaves data nodes to obtain the invention as specified in claims 13, 27 and 41.

48. **Claims 4-5, 18-19 and 32-33** are rejected under 35 U.S.C. 103(a) as being unpatentable over Michael ("High Performance Dynamic Lock-Free Hash Tables and List-Based Sets") in view of AAPA and further in view of Schimmel (US 5,960,434) as

Art Unit: 2187

applied to claims 1-3, 9-11, 13, 15-17, 23-25, 27, 29-31 and 37-39 and 41 above, and further in view of Rakity et al. (US 5,671,446).

49. With respect to **claims 4, 18 and 32**, Michael in view of the AAPA and Schimmel disclose the method (which can also be implemented in a computer-readable storage medium, or by an apparatus) of claims 3, 17 and 31 as shown supra. Michael in view of the AAPA and Schimmel does not disclose the limitation wherein deleting the data node from the linked list additionally involves redirecting the next pointer of the data node to become a back pointer that points to the predecessor of the data node.

50. However, Rakity et al. disclose the limitation deleting the data node from the linked list additionally involves redirecting the next pointer of the data node to become a back pointer that points to the predecessor of the data node (Fig. 8; column 12, lines 46-51).

51. Michael, AAPA, Schimmel and Rakity et al. are analogous art because they are from the same field of endeavor, namely linked list arrangement.

52. At the time of invention, it would have been obvious to a person of ordinary skill in the art to combine the pointer reversal of Rakity et al. with the hash table of Michael and the AAPA.

53. The motivation for doing so would have been so that "the entire LIFO linked list 44 does not have to be scanned every time an element is to be dequeued" (column 12, lines 49-51).

Art Unit: 2187

54. Therefore, it would have been obvious to a person of ordinary skill in the art to combine Rakity et al. with Michael, the AAPA and Schimmel for the benefit of a hash table that uses garbage collection to obtain the invention as specified in claims 4, 18 and 32.

55. With respect to **claims 5, 19 and 33**, Michael in view of the AAPA, Schimmel and Rakity et al. disclose the method (which can also be implemented in a computer-readable storage medium, or by an apparatus) of claims 4, 18 and 32 as shown supra. Michael in view of the AAPA do not disclose the limitations wherein if a search through a chain of nodes from the back pointer does not lead to a live node, the method further comprises:

- gg. obtaining a parent bucket pointer for the bucket pointer;
- hh. searching through the linked list from a node pointed to by the parent bucket pointer to locate a starting node for the bucket pointer; and
- ii. updating the bucket pointer to point to the starting node.

56. However, Schimmel discloses the limitations if a search through a chain of nodes from the back pointer does not lead to a live node (column 8, lines 13-14, if the bucket contains a regular pointer), the method further comprises:

- jj. obtaining a parent bucket pointer for the bucket pointer (column 8, lines 11-16; "if a bucket... includes a logical back pointer," lines 13-14, indicates that the bucket pointer is being initialized);

kk. searching through the linked list from a node pointed to by the parent bucket pointer to locate a starting node for the bucket pointer (column 8, lines 16-20); and

ll. updating the bucket pointer to point to the starting node (column 8, lines 20-22)

57. Michael, AAPA, Schimmel and Rakity et al. are analogous art because they are from the same field of endeavor, namely linked list arrangement.

58. At the time of invention, it would have been obvious to a person of ordinary skill in the art to combine the initialization of bucket pointers of Schimmel with the hash table of Michael, the AAPA and Rakity et al.

59. The motivation for doing so would have been because "re-hashing of data records 156 in step 818 can take too long to be practical" (column 7, lines 39-41).

60. Therefore, it would have been obvious to a person of ordinary skill in the art to combine Schimmel with Michael, the AAPA and Rakity et al. for the benefit of a hash table that initializes new bucket pointers to obtain the invention as specified in claims 5, 19 and 33.

61. **Claims 6, 20 and 34** are rejected under 35 U.S.C. 103(a) as being unpatentable over Michael ("High Performance Dynamic Lock-Free Hash Tables and List-Based Sets") in view of AAPA and Schimmel (US 5,960,434) as applied to claims 1-3, 9-11, 13, 15-17, 23-25, 27, 29-31 and 37-39 and 41 above, and further in view of Hills (US 6,654,773).

Art Unit: 2187

62. With respect to **claims 6, 20 and 34**, Michael in view of the AAPA and Schimmel disclose the method (which can also be implemented in a computer-readable storage medium, or by an apparatus) of claims 2, 16 and 30 as shown supra. Michael in view of the AAPA and Schimmel does not disclose the limitation wherein deleting the data node from the linked list involves using garbage collection or a solution to the repeat offender problem to reclaim the data node if possible.

63. However, Hills discloses the limitation wherein deleting the data node from the linked list involves using garbage collection or a solution to the repeat offender problem to reclaim the data node if possible (column 2, lines 30-48).

64. Michael, AAPA, Schimmel and Hills are analogous art because they are from the same field of endeavor, namely hash table design.

65. At the time of invention, it would have been obvious to a person of ordinary skill in the art to combine the garbage collection method of Hills with the hash table of Michael, the AAPA and Schimmel.

66. The motivation for doing so would have been because "efficient and deterministic garbage collection is achieved" using this particular method (column 2, lines 52-53).

67. Therefore, it would have been obvious to a person of ordinary skill in the art to combine the Hills with Michael, the AAPA and Schimmel for the benefit of a hash table that uses garbage collection to obtain the invention as specified in claims 6, 20 and 34.

68. **Claims 7, 21 and 35** are rejected under 35 U.S.C. 103(a) as being unpatentable over Michael ("High Performance Dynamic Lock-Free Hash Tables and List-Based

Art Unit: 2187

Sets") in view of AAPA and Schimmel (US 5,960,434) as applied to claims 1-3, 9-11, 13, 15-17, 23-25, 27, 29-31 and 37-39 and 41 above, and further in view of Blaker et al. (US 2001/0042204).

69. With respect to **claims 7, 21 and 35**, Michael in view of the AAPA and Schimmel disclose the method (which can also be implemented in a computer-readable storage medium, or by an apparatus) of claims 1, 15 and 29 as shown supra. Michael in view of the AAPA and Schimmel do not disclose the limitation further comprising generating the hash key by performing a pre-hashing operation to achieve a uniform distribution of hash keys over possible hash key values.

70. However, Blaker et al. disclose the limitation further comprising generating the hash key by performing a pre-hashing operation to achieve a uniform distribution of hash keys over possible hash key values [0066, lines 16-17].

71. Michael, AAPA, Schimmel and Blaker et al. are analogous art because they are from the same field of endeavor, namely hash table design.

72. At the time of invention, it would have been obvious to a person of ordinary skill in the art to combine the generation of hash keys with a uniform distribution of Blaker et al. with the hash table of Michael, the AAPA and Schimmel.

73. The motivation for doing so would have been because "if the hash keys which are generated have a random distribution within the data structure address space, then the lower the ratio of entries to table size, the smaller the probability of a 'cluster' of entries of a specific size being created" [0066, lines 7-11].

74. Therefore, it would have been obvious to a person of ordinary skill in the art to combine the Blaker et al. with Michael, the AAPA and Schimmel for the benefit of a hash table that generates a uniform distribution of hash keys to obtain the invention as specified in claims 7, 21 and 35.

75. **Claims 14, 28 and 42** are rejected under 35 U.S.C. 103(a) as being unpatentable over Michael ("High Performance Dynamic Lock-Free Hash Tables and List-Based Sets") in view of AAPA and Schimmel (US 5,960,434) as applied to claims 1-3, 9-11, 13, 15-17, 23-25, 27, 29-31 and 37-39 and 41 above, and further in view of Douceur (US 6,067,547).

76. With respect to **claims 14, 28 and 42**, Michael in view of the AAPA and Schimmel disclose the method (which can also be implemented in a computer-readable storage medium, or by an apparatus) of claims 2, 16 and 30 as shown supra. Michael in view of the AAPA and Schimmel do not disclose the limitations wherein the method further comprises:

- mm. reducing the number of buckets in the bucket array to form a smaller bucket array; and
- nn. using fewer bits from the hash key to perform lookups in the smaller bucket array.

77. However, Douceur discloses the limitations wherein the method further comprises:

oo. reducing the number of buckets in the bucket array to form a smaller bucket array (column 5, lines 24-32); and

pp. using fewer bits from the hash key to perform lookups in the smaller bucket array (column 17, lines 21-34; when searching through a list of nodes, if the node with x significant bits from the key is not found, then a node with $x+1$ significant bits of the key is searched for; when the number of hash buckets is reduced, the number of significant bits of the key used therefore changes from $x+1$ to x).

78. Michael, AAPA, Schimmel and Douceur are analogous art because they are from the same field of endeavor, namely hash table design.

79. At the time of invention, it would have been obvious to a person of ordinary skill in the art to combine the dynamic subtraction of hash table buckets of Douceur with the hash table of Michael and the AAPA.

80. The motivation for doing so would have been to "free memory for use by other applications" by reducing the amount of memory used by the hash table (column 5, lines 25-26).

81. Therefore, it would have been obvious to a person of ordinary skill in the art to combine the Douceur with Michael, the AAPA and Schimmel for the benefit of a hash table that dynamically subtracts hash buckets from itself to obtain the invention as specified in claims 14, 28 and 42.

Response to Arguments

82. Applicant's arguments filed 9/18/2006 have been fully considered but they are not persuasive.

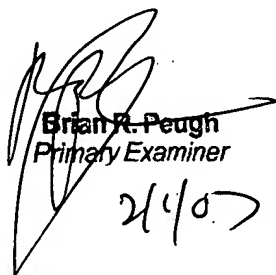
83. In the second paragraph on page 13 through the first paragraph of page 14 of Applicant's arguments filed 9/18/2006, with respect to the rejection of claims 1, 15, and 29 under 35 USC 103(a), Applicant argues "...*the combined system of Michael, AAPA, and Schimmel teaches **adding a pointer to the new buckets which point to the old buckets**, and then rehashing the data records over time*" "In contrast, the present invention **directly copies the old buckets into the top half of the new bucket array** (see FIG. 2C and paragraphs [0043]-[0049] of the instant application). This is beneficial because it does not require the subsequent rehashing of the data records. There is nothing within Michael, AAPA, or Schimmel, either separately or in concert, which suggests directly copying the old buckets into the top half of the new bucket array." (emphasis in original) The Examiner respectfully disagrees. While Schimmel does teach that the new buckets contain pointers to the old buckets which point to the old hash buckets, as shown in figure 5 A and discussed at column 7 lines 56-64, only the new buckets 410-416 contain the logical back pointers. Buckets 320-326 are the old buckets. This is further illustrated at column 7 lines 44-46, which shows that hash table 310, which is shown in figure 3A to contain four buckets, is dynamically resized as hash table 508 of figure 5A. Accordingly, buckets 320-326 are copied from hash table 310 to hash table 508.

Conclusion

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Jared I. Rutz whose telephone number is (571) 272-5535. The examiner can normally be reached on M-F 8:00 AM - 4:00 PM.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Donald Sparks can be reached on (571) 272-4201. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.


Brian R. Peugh
Primary Examiner
2/1/07

Jared I Rutz
Examiner
Art Unit 2187

jir
2/1/07